

REMARKS

In response to the Office Action mailed June 23, 2004, Applicants respectfully request reconsideration. To further the prosecution of this Application, Applicants submit the following remarks, have canceled claims and have added new claims. The claims as now presented are believed to be in allowable condition.

Claims 1-23, 25, 30, 32 and 34-42 were pending in this Application. By this Amendment, claims 43-50 have been added. Accordingly, claims 1-23, 25, 30, 32 and 34-50 are now pending in this Application. Claims 1, 14, 23 and 32 are independent claims.

Preliminary Matters

First, Applicants wish to point out that Applicants submitted an IDS on August 12, 2004 after the Patent Office mailed the current Office Action. Applicants respectfully request that the Examiner complete the PTO-1449 contained within that IDS and return that PTO-1449 form with the next correspondence from the Patent Office.

Second, Applicants wish to point out that Applicant have amended the Attorney Docket Number. The new Attorney Docket Number is EMC04-49(01195).

Third, Applicants wish to point out that all future correspondences should be sent to Applicants' Representative at the address below:

David E. Huang, Esq.
Attorney for Applicants
Registration No.: 39,229
CHAPIN & HUANG, L.L.C.
Westborough Office Park
1700 West Park Drive
Westborough, Massachusetts 01581
Telephone: (508) 366-9600
Facsimile: (508) 616-9805

Applicants will be submitting a POWER OF ATTORNEY TO PROSECUTE APPLICATIONS BEFORE THE USPTO together with a statement under 37 CFR 3.73(b) (Form PTO/SB/96 or equivalent) in compliance with 37 CFR 1.31 and 1.33 shortly.

Objection to the Drawings

The drawings were objected to as not showing every feature of the invention as specified in the claims.

Applicants have amended the claims to remove the claim language in question. Accordingly, the objection to the drawings should be withdrawn.

Provisional Obviousness-Type Double Patenting Rejection

Claims 4, 5 and 23 were provisionally rejected under the judicially created doctrine of obviousness-type double patenting in view of claims 1, 4 and 20, respectively, of co-pending U.S. Patent Application No. 10/004,090.

Applicants wish to point out that the provision rejection has been removed in co-pending U.S. Patent Application No. 10/004,090. If a Terminal Disclaimer is still required in the current Application, Applicants will delay submission of such a Terminal Disclaimer in compliance with 37 C.F.R. 1.321(c) until the other claim rejections are removed.

Rejections under §102 and §103

Claims 32-37 were rejected under 35 U.S.C. §102(b) as being anticipated by U.S. Patent No. 5,805,787 (Brant et al.). Claims 1-23, 25, 30, and 38-42 were rejected under 35 U.S.C. §103(a) as being unpatentable over Brant et al. in view of U.S. Patent No. 5,195,100 (Katz et al.).

Applicants respectfully traverse the rejection of claims 32-37 and request reconsideration of these claims. Additionally, Applicants submit that amended claims 1-23, 25, 30, and 38-42 patentably distinguish over the cited prior art and request reconsideration of these claims as well. The claims are now in allowable condition.

Brant presents a hierarchy of various contemporary storage configurations in order of highest cost but fastest performance first and lowest cost but slowest performance last (column 5, lines 11-15). For example, a disk based disk cache is listed as being lower in the hierarchy vis-à-vis registers of a microprocessor, but

higher in the hierarchy vis-à-vis optical, tape and/or library storage (column 5, lines 16-28). In connection with the hierarchy of contemporary storage configurations, Brant mentions certain RAID approaches, e.g., that RAID3 is supported with daughter cards with flow through parity generators (column 5, lines 29-45).

In connection with the Brant invention, Brant discloses a host 10 in bidirectional communication with a controller 20 of a disk based disk cache subsystem 16 via a link 11 (column 5, lines 51-53 and Fig. 1). The controller 20 employs its own buffer to interface with a separate disk based disk cache array 22 (column 5, lines 54-56). Brant then proposes use of a direct mapped strategy to map logical blocks to a physical location where the direct mapped strategy always writes a specific logical block to the same physical location so that the search entails a constant time segment (column 6, line 59 through column 7, line 5). A typical direct mapped write operation writes N copies of the data to the disk based disk cache comprising of N disks with one copy to each disk (column 7, lines 6-8). To select the device which can most quickly access data (column 7, lines 15-16), devices are preferably synchronized so that the rotation of each device is offset by $1/N$ (where N is the number of devices in the system) to be able (column 7, lines 23-25). Brant further discloses, as other mapping strategies and among other things, a set associative strategy (column 8, lines 42-55), a fully associative LRU cache management strategy (column 8, line 56 through column 9, line 5), and a fully associative hybrid LRU/fastest fit strategy (column 9, lines 6-12).

Katz discloses methods and apparatus for detecting and correcting various data errors (Abstract). Katz points out that, in certain cases, there may not be sufficient time to complete an operation in the event of a power failure (column 12, lines 14-16). Therefore, before any write operation is started on any disk, within a non-volatile memory 413 is stored a journal of information concerning a CPU write request and the write operations to be performed (column 12, lines 17-21).

Claims 1-13, 25 and 39

Claim 1, as amended, is directed to a cache which includes a front-end interface that receives data access requests that specify respective data storage addresses, a back-end interface that can retrieve data identified by the data storage addresses, cache storage formed by at least two disks, and a cache manager that services at least some of the requests received at the front-end interface using data stored in the cache storage. The cache manager is configured to receive a write request to store data and, in response to the write request, split that data into data portions and separately store the data portions on respective disks of the cache storage. The cache manager is further configured to receive a read request to read the data and, in response to the read request, concurrently read the data portions which are separately stored on the respective disks of the cache storage to retrieve the data.

The cited prior art does not teach or suggest, either alone or in combination, a cache having a cache manager which is configured to (i) receive a write request to store data and, in response to the write request, split that data into data portions and separately store the data portions on respective disks of the cache storage, and (ii) receive a read request to read the data and, in response to the read request, concurrently read the data portions which are separately stored on the respective disks of the cache storage to retrieve the data, as recited in claim 1. Rather, Brant discloses a controller 20 and a separate disk based disk cache array 22 (e.g., see column 5, lines 54-56 of Brant), where a typical direct mapped write operation writes N copies of the data to the disk based disk cache comprising of N disks with one copy to each disk (e.g., see column 7, lines 6-8 of Brant). Since Brant writes one copy of the data to each disk, Brant does not disclose a cache manager configured to split data into data portions and separately store the data portions on respective disks of a cache storage, as recited in claim 1.

Moreover, to select the device which can most quickly access data (e.g., see column 7, lines 15-16 of Brant), Brant discloses synchronizing devices so that the rotation of each device is offset by 1/N (where N is the number of devices in the

system) (e.g., see column 7, lines 23-25 of Brant). In view of such synchronization, it is unclear how one could modify Brant to have a cache manager configured to split data into data portions and separately store the data portions on respective disks of a cache storage, as recited in claim 1. If anything, it seems that splitting data into data portions and separately storing the data portions on respective disks would prevent Brant from being able to select the device which can most quickly access data. Katz which was cited by the Office Action on page 14, last paragraph, for disclosing storage of a journal of information concerning a CPU write request and the write operations to be performed in case of a power failure (e.g., see column 12, lines 17-21 of Katz) does not teach or suggest how one could modify Brant to split data into data portions and separately store the data portions on respective disks, as required by claim 1, while still being able to select the device which can most quickly access data as disclosed in Brant.

For the reasons stated above, claim 1 patentably distinguishes over the cited prior art, and the rejection of claim 1 under 35 U.S.C. §103(a) should be withdrawn. Accordingly, claim 1 is now in allowable condition.

Because claims 2-13, 25 and 39 depend from and further limit claim 1, claims 2-13, 25 and 39 are in allowable condition for at least the same reasons.

Claims 14-22 and 40

Claim 14, as amended, is directed to a method of servicing data access requests at a cache. The method includes receiving the data access requests at the cache. The cache has cache storage formed by at least two disks. The requests specifying respective data storage addresses. The method further includes servicing at least some of the requests using data stored in the disks. Receiving the data access requests includes receiving a write request to store data, and servicing at least some of the requests includes (i) splitting that data into data portions and (ii) separately storing the data portions on respective disks of the cache storage in response to the write request. Additionally, receiving the data access requests further includes receiving a read request to read the data, and servicing at least

some of the requests further includes concurrently reading the data portions which are separately stored on the respective disks of the cache storage to retrieve the data in response to the read request.

The cited prior art does not teach or suggest receiving a write request to store data, and (i) splitting that data into data portions and (ii) separately storing the data portions on respective disks of a cache storage in response to the write request, as recited in claim 14. In contrast, as mentioned above in connection with claim 1, Brant discloses an approach which involves writing one copy of the data to each disk. Furthermore, there is no teaching or suggestion in either Brant or Katz as to how one could modify the Brant approach to split data into data portions and separately store the data portions on respective disks of a cache storage, as recited in claim 14.

Accordingly, claim 14 patentably distinguishes over the cited prior art for at least the same reasons as claim 1. Thus, the rejection of claim 14 under 35 U.S.C. §103(a) should be withdrawn and claim 14 is now in allowable condition.

Because claims 15-22 and 40 depend from and further limit claim 14, claims 15-22 and 40 are in allowable condition for at least the same reasons.

Claims 23, 30 and 41

Claim 23, as amended, is directed to a data storage system which includes a back-end storage system having an address space. Addresses in the address space identify blocks of storage. The data storage system further includes a cache, for the back-end storage system, having a lesser storage capacity than the back-end storage system. The cache includes a front-end interface that receives I/O (Input/Output) requests that specify respective addresses of back-end storage blocks, a back-end interface that communicates with the back-end storage system, cache storage formed by at least two disks having platter diameters less than 3.5 inches, and a cache manager that services at least some of the I/O requests received via the front-end interface using blocks temporarily stored in the cache storage. The cache manager is configured to receive a write request to store data

and, in response to the write request, split that data into data portions and separately store the data portions on respective disks of the cache storage. The cache manager is further configured to receive a read request to read the data and, in response to the read request, concurrently read the data portions which are separately stored on the respective disks of the cache storage to retrieve the data.

The cited prior art does not teach or suggest a cache manager which is configured to receive a write request to store data and, in response to the write request, split that data into data portions and separately store the data portions on respective disks of a cache storage, as recited in claim 23. Rather, as mentioned above in connection with claim 1, Brant discloses an approach which involves writing one copy of the data to each disk. Moreover, there is no teaching or suggestion in either Brant or Katz as to how one could modify the Brant approach to split data into data portions and separately store the data portions on respective disks of a cache storage, as recited in claim 23.

Accordingly, claim 23 patentably distinguishes over the cited prior art for at least the same reasons as claim 1. Therefore, the rejection of claim 23 under 35 U.S.C. §103(a) should be withdrawn and claim 23 is now in allowable condition.

Because claims 30 and 41 depend from and further limit claim 23, claims 30 and 41 are in allowable condition for at least the same reasons.

Claims 32, 34-38 and 42

Claim 32 is directed to a data storage system which includes a back-end storage system having a back-end address space, addresses in the address space identifying blocks of storage. The data storage system further includes a plurality of caches for the back-end storage system. Each of the plurality of caches has a lesser storage capacity than the back-end storage system. Each of the plurality of caches includes a front-end interface that receives I/O (Input/Output) requests that specify respective addresses of back-end storage blocks, a back-end interface capable of communicating with one of back-end storage system and another of one of said plurality of caches, and cache storage formed by at least two disks. The

cache storage has a respective cache storage address space. The plurality of caches further includes a cache manager that services at least some of the I/O requests received via the front-end interface using blocks temporarily stored in the data storage system. At least some of the I/O requests correspond to addresses in the respective cache storage address space of at least some of the plurality of caches.

Brant does not teach a data storage system which includes, among other things, a plurality of caches for a back-end storage system where each of the plurality of caches has a lesser storage capacity than the back-end storage system, and where each of the plurality of caches includes (i) a front-end interface that receives I/O requests that specify respective addresses of back-end storage blocks, (ii) a back-end interface capable of communicating with one of back-end storage system and another of one of said plurality of caches, (iii) cache storage formed by at least two disks, and (iv) a cache manager that services at least some of the I/O requests received via the front-end interface using blocks temporarily stored in the data storage system, as recited in claim 32. Rather, Brant discloses a host 10 in bidirectional communication with a controller 20 of a disk based disk cache subsystem 16 via a link 11 (e.g., see column 5, lines 51-53 of Brant). There is no other plurality of caches between the Brant host 10 and the disk system 25, as recited in claim 32 (e.g., see Fig. 1 of Brant).

In the Office Action's attempt to establish the rejection of claim 32, the Office Action contends that "A plurality of caches is disclosed in column 5, lines 29-31, which discusses that the invention of Brant et al may be put in any multiple locations in the hierarchy of the system," (see the first sentence of page 9 of the Office Action). Applicants respectfully submit that the Office Action has misconstrued the teachings of Brant. The cited portion of Brant does not disclose that the invention of Brant may be put in any multiple locations in a hierarchy of a data storage system, as the Office Action contends. Rather, the cited portion of Brant merely presents a hierarchy of various contemporary storage configurations in order of highest cost but fastest performance first and lowest cost but slowest performance last for

comparison purposes (e.g., see column 5, lines 11-15 of Brant). For example, a disk based disk cache is listed as being lower in the hierarchy vis-à-vis registers of a microprocessor, but higher in the hierarchy vis-à-vis optical, tape and/or library storage (e.g., see column 5, lines 16-28 of Brant). Accordingly, there is no disclosure of a data storage system as recited in claim 32 by Brant, but instead, the Office Action has inadvertently taken a statement by Brant out of context.

"A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference."¹ "The identical invention must be shown in as complete detail as is contained in the ... claim."² Since Brant does not describe each and every element as set forth in claim 32, claim 32 is not anticipated by Brant.

For the reasons stated above, claim 32 patentably distinguishes over the cited prior art, and the rejection of claim 32 under 35 U.S.C. §102(b) should be withdrawn. Accordingly, claim 32 is in allowable condition.

Because claims 34-38 and 42 depend from and further limit claim 32, claims 34-38 and 42 are in allowable condition for at least the same reasons.

Newly Added Claims

Claims 43-51 have been added and are believed to be in allowable condition. Claims 43-44 depend from claim 1. Claims 45-46 depend from claim 14. Claims 47-48 depend from claim 23. Claims 49-51 depend from claim 32. Support for claims 43, 45, 47 and 50 is provided within the Specification, for example, on page 8, lines 15-21 and Fig. 1 which shows extending connections for disks. Support for claims 44, 46, and 48 is provided within the Specification, for example, on page 10, lines 9-22 and Fig. 8. Support for claim 49 and the amendments to claims 1, 14 and 23 is provided within the Specification, for example, on page 9, last line through page 10, line 8 and Figs. 5-7. No new matter has been added.

¹ *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631, 2 USPQ2d 1051, 1053 (Fed. Cir. 1987).

² *Richardson v. Suzuki Motor Co.*, 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989).

-21-

Conclusion

In view of the foregoing remarks, this Application should be in condition for allowance. A Notice to this affect is respectfully requested. If the Examiner believes, after this Amendment, that the Application is not in condition for allowance, the Examiner is respectfully requested to call the Applicants' Representative at the number below.

Applicants hereby petition for any extension of time which is required to maintain the pendency of this case. If there is a fee occasioned by this Amendment, including an extension fee, that is not covered by an enclosed check, please charge any deficiency to Deposit Account No. 50-0901.

If the enclosed papers or fees are considered incomplete, the Patent Office is respectfully requested to contact the undersigned collect at (508) 366-9600, in Westborough, Massachusetts.

Respectfully submitted,



David E. Huang, Esq.
Attorney for Applicants
Registration No.: 39,229
CHAPIN & HUANG, L.L.C.
Westborough Office Park
1700 West Park Drive
Westborough, Massachusetts 01581
Telephone: (508) 366-9600
Facsimile: (508) 616-9805

New Attorney Docket No.: EMC04-49(01195)

Dated: October 25, 2004